The Dynamics of Information Fusion: Synthetic Versus Misassociation

Kessel, Ronald T.

August 2006

Report Documentation Page

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE **AUG 2006**	2. REPORT TYPE	3. DATES COVERED **00-08-2006 to 00-08-2006**

4. TITLE AND SUBTITLE **The Dynamics of Information Fusion: Synthetic versus Misassociation**	5a. CONTRACT NUMBER
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **NATO Undersea Research Centre,Viale San Bartolomeo 400,19138 La Spezia, Italy, ,**	8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES
The original document contains color images.

14. ABSTRACT

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**		**9**	

NATO Undersea Research Centre (NURC)

NURC conducts world class maritime research in support of NATO's operational and transformational requirements. Reporting to the Supreme Allied Commander Transformation, the Centre maintains extensive partnering to expand its research output, promote maritime innovation and foster more rapid implementation of research products.

The Scientific Programme of Work (SPOW) is the core of the Centre's activities and is organized into four Research Thrust Areas:

- Expeditionary Mine Countermeasures (MCM) and Port Protection (EMP)
- Reconnaissance, Surveillance and Undersea Networks (RSN)
- Expeditionary Operations Support (EOS)
- Command and Operational Support (COS)

NURC also provides services to other sponsors through the Supplementary Work Program (SWP). These activities are undertaken to accelerate implementation of new military capabilities for NATO and the Nations, to provide assistance to the Nations, and to ensure that the Centre's maritime capabilities are sustained in a fully productive and economic manner. Examples of supplementary work include ship chartering, military experimentation, collaborative work with or services to Nations and industry.

NURC's plans and operations are extensively and regularly reviewed by outside bodies including peer review of the research, independent national expert oversight, review of proposed deliverables by military user authorities, and independent business process certification. The Scientific Committee of National Representatives, membership of which is open to all NATO nations, provides scientific guidance to the Centre and the Supreme Allied Commander Transformation.

NOTE: The NURC Reprint series reprints papers and articles published by NURC authors in the open literature as an effort to widely disseminate NURC products. Users are encouraged to cite the original article where possible.

The Dynamics of Information Fusion: Synthesis Versus Missassociation

Ronald T. Kessel

NATO Undersea Research Centre,
La Spezia, Italy
Kessel@nurc.nato.int

3 Feb 2007

Abstract - *Fusion entails both the association and synthesis of information. If missassociations occur, they obviously undermine the gains won by synthesis, compromising the fusion product. An analytic framework is presented here to study the competition between the negative effect of missassociation and the positive effect of synthesis, to demonstrate and analyze their interplay quantitatively. Here the quality of information being enhanced or degraded is taken to be the extent to which the information correctly determines a decision or action inference. To say that the uncertainty injected by missassociation may overwhelm the uncertainty reduction won by synthesis, for instance, would mean that this inference-determining quality of information falls in fusion below that of the best information source working independently. This is ultimately a study in uncertainty dynamics: the beneficial reduction of uncertainty by synthesis in fusion versus the detrimental increase of uncertainty due to association, which are both always present in fusion to some degree.*

Keywords: Data association, missassociation, information quality, fusion performance, fusion theory.

1 Introduction

Fusion can be represented schematically as in Fig. (1). The arrows indicate the flow of information, from various sources coming together to produce another information product on the right, the output presumably more useful in some important respect. The fusion occurs inside the box, which could represent the human mind, an expert system, a combined human-machine system, and so forth. Beginning from this simple schematic, we might ask if there are principles governing the dynamics of the *quality* of information in fusion. Is it true, for instance, that fusion necessarily *improves* maters? *Are there conditions, for instance, under which improvement might for some reason be impossible?*

Fusion entails association and synthesis—the *association* first of all of subject matter between the information sources to be fused, and then the *synthesis* of the information pertaining to that subject matter [1]. Synthesis is often the focus of discussion in fusion. It
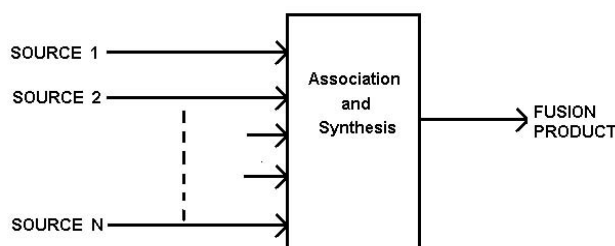

Figure 1: A conceptual schematic of fusion. Association is the correlation of subject matter treated by all incoming information sources. Synthesis is the combination of that information into a single fusion product.

is what adds value to information. If it were a matter of synthesis alone, then the prospects for fusion (in theory at least) are always good. More independent information generally means better information in total. What must not be forgotten, however, is that is there can be no fusion without association.

Illustrating briefly, imagine a surveillance aircraft on a routine mission overflying and photographing a ship at sea. The crew on board the aircraft hail the ship to collect routine information about destination, purpose, and so forth, which in this case apparently contradicts the heading of the ship and the equipment evident on its decks. The crew in the air grow suspicious, alert surface forces for interception, and so forth, not realizing, however, that their visual contact was with one ship, while their radio contact was with *another* very distant ship. The fusion between independent visual and radio information took place in the mind of the crew on the implicit assumption that there was just one subject, whereas, in reality there were two. It is a simple error of association. Elsewhere a single subject might be treated as several. When associating information between several databases for instance, using the name of a person or a ship to relate records, the name may be entered differently in any number of places, in any number of ways, by typing errors, changes in convention, omissions, which make the one subject appear as several independent subjects.

There are of course ways to avoid these kinds of association errors. The points to be made nonetheless

are 1) that when dealing with situations of increasing complexity, the uncertainty about correct association, or the probability of missassociation, tends as a rule to increase; and 2) that the more independent the incoming sources of information are, the more potential they hold for gains through synthesis, but the more uncertainly they are also associated together. Thus the system developer is likely at some stage to face diminishing returns. Fusion may actually fail to improve matters, and may even significantly degrade matters. Much as synthesis holds only positive potential of improving information, association holds only negative potential for degrading it, and both are *always* at work to some degree in fusion. Our question for fusion may therefore be restated this way: *To what extent will the probabilities of missassociation undermine the improvements by synthesis?*

This paper addresses the question analytically, in a generic way, as a step (it is hoped) toward developing engineering rules of thumb by which to judge the prospects for fusion, by analysis, before a fusion system is built and fielded, and to help identify and diagnose problems with fusion in operation.

That fusion entails association is of course nothing new. But in the analysis of fusion it is often either explicitly asserted or implicitly assumed that the subject matter treated by all of the incoming sources of information can be unambiguously associated; that two independent subjects have not mistakenly been confused as one, or that one subject has not been confused as two or more. What has been missing is the admission that association errors can occur in practice, with a subsequent quantitative analysis of what their affect is expected to be. In some cases association uncertainties have been ignored because the application is of such a tightly controlled or consistent nature that an unambiguous correspondence between the subject matter treated by each information source was assumed to be assured by design. This is generally (though not strictly) true of the biological examples often cited in the early fusion literature, such as the fusion of sound and sight in animals hunting prey; the multiple sensors in these cases being co-located, synchronous, and correlated by a brain evolved for the purpose. Association uncertainties have also been ignored because the information sources were assumed (often mistakenly in practice) to be totally error and ambiguity free. The information in data bases, for instance, such as in a daily report of harbour vessel traffic and the historic record of vessel traffic, might treated as if uncontaminated by entry errors, spelling mistakes, omissions or duplications, synonymic variations, ontological ambiguities, and so forth, whereas, what was one day entered as a "fishing boat" might on other days been entered as a "trawler", ship names and codes may have been entered with spelling errors, and so forth, all of which undermines associations made between data bases, and even within the same data base. Elsewhere association uncertainties might be ignored if their effects would be so obviously dangerous that operation under any shadow of association uncertainty

would be unthinkable. One either operates with full confidence in association, or one does not operate at all.

1.1 Scope

A conceptual model of fusion is proposed here that illustrates the competition of forces quantitatively, for a particular kind of fusion, when the fusion product is intended to determine a particular *decision/action inference* with more reliability than any of the fused sources of information working independently could determine the same inference. This would include, for instance, the fusion of various target detection and classification sensors, the fusion of decisions of various experts (human or expert system) on one given question, or the fusion of various data bases to detect anomalous activity for counter terrorism or maritime security. this would not include fusion for (continuous) parameter estimation, such as the refinement by multi-sensor fusion of a target's range and bearing, or fusion that merely increases coverage, such as the side-by-side presentation of information from sensors viewing independent but neighboring regions of the earth.

By decision/action inference is meant that the purpose of information gathering and fusion is to support making decisions or taking action. It may be the decision to change one's stance or attitude toward a radar contact, to commit resources to further information gathering, to continue with one's present actions without change, to issue a command or raise an alarm, and so forth. In every case an agent, human or machine, considers information in light of some decision to be made, or action to be taken, that is intimately related to that agent's situation awareness. The utility of the information then depends above all else on the reliability with which the inference is made. It also depends very strongly on the nature of the application. The analysis therefore begins by relating utility measures to the performance probabilities for inference, with the performance probabilities serving as generic surrogate metrics for more specialized utilities. Unlike information theory, then, we are asking about the utility of information, not about its quantity in bits, which correlates poorly with its ultimate utility in human terms. Information is good or bad insofar as it produces good or bad decisions or actions. The question posed earlier for fusion may therefore be restated more precisely again: *To what extent will a probability of missassociation in fused information undermine the decision/action inferences based on that information?*

For simplicity in the final model, the decision/action inference is assumed to be, in effect, a choice between just two predefined options, labeled 0 and 1, as in target detection and classification, for example, when an object in a scene that is under surveillance by several sensors is classified as a target (1) or clutter (0). But any two-option inference is possible. And just two sources of information are considered for fusion. The approach is generic inasmuch as the way in which decision/action inferences are made, by human or automation, and the way in which, or algorithm by which,

the fusion is carried out, are for the most part of no consequence. Performance and utility are the key parameters.

2   Utility of information

The quality of information lies ultimately in the quality of the decision/action inference that might be made from it, which is the probability that the inference is correctly made. This is the conditional probability P^C that, given a world in state m, the appropriate resource-affecting inference m is made about that state of the world, on the basis of the information available amid uncertainty,

$$P_m^C = P\left(\text{ infer } m \mid \text{given state } m\right). \qquad (1)$$

P_m^C is not necessarily a measure of utility itself, but it is a driver behind any utility metric U when being correct in decision and action is paramount. Whatever U may be, that is to say, it will be a monotonic function of the conditional probabilities P_m^C,

$$U = U\left(P_m^C\right) \quad \text{and} \quad \frac{\partial U}{\partial P_m^C} \geq 0 \text{ for all } m, \qquad (2)$$

assuming that higher U means higher utility. U might be the probability of being correct during a mission

$$U = \sum_{m=0}^{M-1} p_m P_m^C, \qquad (3)$$

for example, in which a possibly relevant state of the world m is expected to appear with relative frequency p_m, and there are M such states. Or U might be the expected benefit (negative expected cost) of operations. There are of course other dimensions to the utility of information, like timeliness, distribution, and access, but here we are interested in the information that happens to be available at the time a decision/action inference is made, not in the information that might have been available.

2.1   Performance probabilities

Performance probabilities for two-option inferences can be modeled very generally as a thresholded decision variable. That is to say, given one of two mutually exclusive states of the world, labeled $m = 0$ or 1, the expected (average) performance of a human or a machine making an inference amid uncertainty, $I = 0$ or 1, on the basis of their information about which state holds, can generally be modeled as if I were determined by the value of a decision variable c relative to a threshold c_T,

$$I = \begin{cases} 0 & \text{if } \quad c < c_T, \\ 1 & \text{if } \quad c \geq c_T, \end{cases} \qquad (4)$$

where c is a random variable drawn from a probability distribution conditioned by the true state of the world, as shown in Fig.(2). This kind of model is used in physchophysics for human behavior[2], for instance,
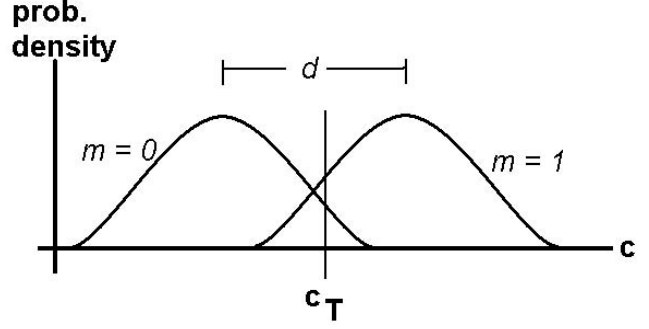

Figure 2: The performance of an inference agent is modeled as a thresholded random decision variable c whose distribution is conditioned by the state of the world, of which there are only two of interest in this case, $m = 0$ and 1. Greater separation d between the distributions means better quality inferences.

without claiming that c or c_T exist in reality, and without making any assertions about the process by which the assertion was made. c and c_T are simply a means for modeling the average (expected) performance or reliability of inferences made amid uncertainty, without predicting or explaining individual inferences. The model in Fig.(2) is used in much the same way here, making no claims about what the information may be, no claims about the process or algorithm by which the decision-action inference is made, and no claims about what a particular individual inference will be. The model therefore applies very generally to many different situations.

In the example below, the density functions in Fig.(2) are for simplicity taken to be Gaussian of unit variance, with separation d between means. Increasing the separation implies that either the information or the inference agent, or both, have been improved. In this paper it is the quality of information that is at issue, particularly as the probability of missassociation is increased in fusion, so the intrinsic ability of any inference agent is assumed to be constant.

Figure (3) shows the range of performance probabilities P_0^C and P_1^C that might result for a given information source as the inference agent's aversion to making one or another type of error changes, which is to say, in terms of the model in Fig.(2), as the threshold c_T is varied along the c axis. Each point $\left(P_0^C, P_1^C\right)$ on a curve in Fig.(3) corresponds to a different value of c_T for a given quality of information. This parametric representation of performance is much the same as a receiver operator characteristic (ROC) curve in detection or classification, but with the horizontal axis reversed; an ROC curve typically having the probability of false alarm P_{FA} plotted along the horizontal axis rather than $P_0^C = 1 - P_{FA}$. A uniform treatment of all P_m^C as in Fig.(3) simplifies the analysis when multi-option inferences ($m = 0, 1, 2 \ldots M - 1$) are considered, when the parametric curve becomes a multi-dimensional surface.
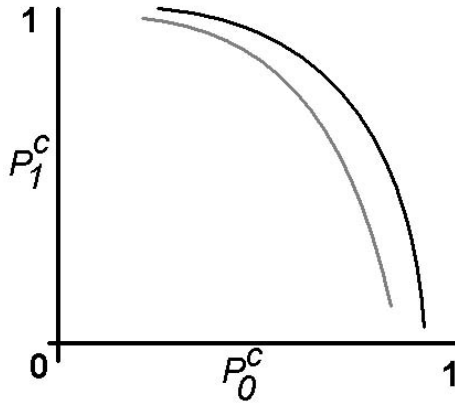
Figure 3: The performance probabilities P_0^C and P_1^C are plotted here as the threshold c_T is varied. The lighter curve signifies an inference of lower utility because for equal P_0^C the darker curve has a higher P_1^C.

Figure 4: The two-dimensional extension of Fig.(2) for the fusion of two sources of information.

2.2   Comparing information quality

Given N information sources labeled $s = 1, 2, \ldots N$, one would have have N different models of the type in Fig.(2), each with its own decision variable, threshold, and probability density functions, and performance curve in Fig.(3). In general, the form of the utility function U would have to be known in order to rank all of these from best to worst in terms of their ultimate utility for the same application. But there are conditions owing to monotonicity (2) under which the best or worst source can be identified, regardless of what U may be. A source $s = t$, for instance, whose performance probabilities $P_{m,t}^C$ (with second subscript for source label) are maximum in the following respect

$$P_{m,t}^C \geq P_{m,s}^C \text{ for all } m \text{ and } s \neq t, \qquad (5)$$

will be the *best* source of information, offering the highest utility in decision or action. On the other hand, a source whose performance probabilities are minimum

$$P_{m,t}^C \leq P_{m,s}^C \text{ for all } m \text{ and } s \neq t \qquad (6)$$

will be the *worst* source of information, offering the lowest utility. These are sufficient, but not necessary, conditions for being best or worst. They apply very generally to any application, allowing us to determine in some cases whether fusion rises to the top, or falls to the bottom rank among the information sources fused.

For simplicity, we continue now with the two-option inference, $m = 0, 1$, and the fusion of just two information sources, $s = 1, 2$, demonstrating fusion when associations are unambiguous, and again with association uncertainty (a probability of missassociation).

3   Fusion with association uncertainty

Fusion generally divides into two types. The simplest, *decision fusion*, is when inferences I_s made on the strength of each information source s independently
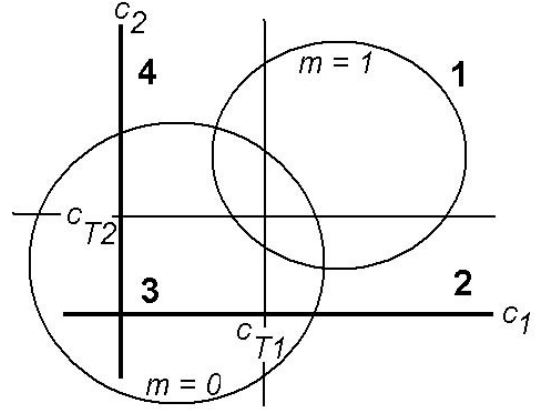
are fused together into a single inference I_F, typically using a logical AND or OR operation, or a weighted voting scheme[3] ,[4]. And there is *feature fusion*, in which the decision variables c_s or other inference-determining features are fused, creating in effect another decision variable c_F by which a single fused inference I_F is made. In principle, feature fusion offers higher utility gains, though often not dramatically high in practice because thorough knowledge of the probability distributions (as in Fig.(2)) is required to optimize performance, which is rarely available practice. Decision fusion, on the other hand, tends to be more robust in practice.

Here we use a simple form of feature fusion for illustration. Assume that each of two sources for fusion promise equal utility $P_{m,1}^C = P_{m,2}^C$, both being derived from Gaussian conditional probability distributions of unit variance as in Fig.(2), with separation between means of $d = 2.0$. Equality in performance probabilities implies equal expected utility U in independent operation, but does not imply equality between inferences (i.e. c_1 need never equal c_2, nor I_1 equal I_2). Assume, moreover, that the information provided by each source is independent. Then the joint conditional distributions of decision variables (c_1, c_2) are as shown in Fig.(4), which is a two-dimensional extension of Fig.(2) for two-source fusion. That is to say, given an occasion for inference, the inference-determining cues in each information source each contribute in effect some net inference-determining cue that has, it can be assumed, roughly the same independent determining power as when the two sources of information are used independently. Indeed, each individual source contributes, in effect, its decision variable c_m on which its independent performance depends. Each occasion for inference therefore generates the ordered (c_1, c_2). A simple, admittedly suboptimal way of modeling the fusion of the two sources into a single inference is again to threshold each decision variable as shown in Fig.(4). If (c_1, c_2) lies in quadrant #1, then the fusion result is $I_F = 1$; otherwise $I_F = 0$. In the absence of association errors this is in fact equivalent to decision fusion with a logical AND rule.
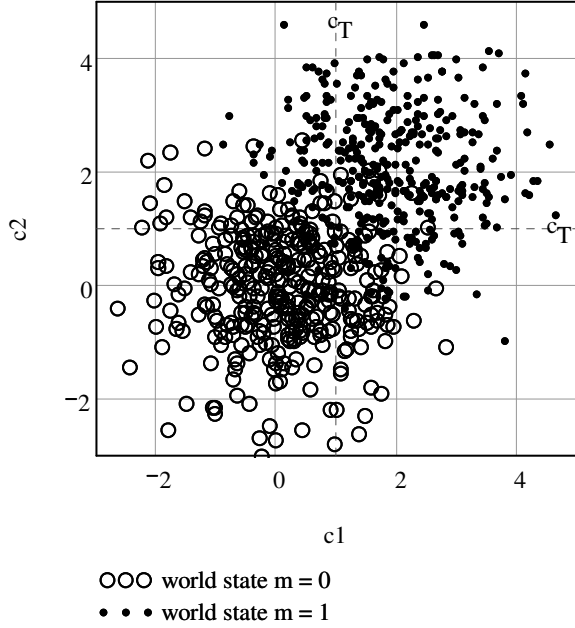
Figure 5: A scatter plot of points for the MonteCarlo simulation. The distributions here are arranged much like those shown schematically in Fig.(xx). 750 points are plotted of the 5000 used in the simulations. The states $m = 0$ and 1 appear with equal frequency.

Figure (5) shows the resulting scatter plot when many occasions for inference are modeled using Monte-Carlo simulation, when the two real-world possibilities $m = 0$ and 1 occur with equal frequency. The performance probabilities are estimated by counting the number of points in various quadrants of the (c_1, c_2) plane created by the thresholds c_T:

$$
\begin{aligned}
\text{source 1} \quad & P_{0,1}^C = & \text{(points in quadrants 3, 4)}/n \\
\text{source 1} \quad & P_{1,1}^C = & \text{(points in quadrants 1, 2)}/n \\
\text{source 2} \quad & P_{0,2}^C = & \text{(points in quadrants 2, 3)}/n \\
\text{source 2} \quad & P_{1,2}^C = & \text{(points in quadrants 1, 4)}/n \\
\text{fusion 1\&2} \quad & P_{0,F}^C = & \text{(points in quadrants 2, 3, 4)}/n \\
\text{fusion 1\&2} \quad & P_{1,F}^C = & \text{(points in quadrant 1)}/n
\end{aligned}
$$

(7)

in which n is the total number of points in the scatter plot. Figure (6) shows the result $P_{m,s}^C$ when $n = 5000$, and simulation is repeated times with different values of threshold c_T. The simple fusion is clearly superior to either source working independently because its performance probabilities are to the right and above that of both sources working independently (the condition (5)). No possibility of association error has been admitted here because the subject matter treated by the two sources were unambiguously associated by design, in the way that the ordered pairs (c_1, c_2) were constructed: one pair for each real-world occasion for inference, and each real-world occasion calling for just one inference.

Now let us introduce the prospect of association errors by assuming that there exists a class of real-world situations in which the occasions for inference come in pairs. This might happen in maritime surveillance, for instance, when a scene contains two ships, A and B. This potentially leads to a confusion in subject matter between the two information sources because the two contacts in one sensor's view may be associated in two different ways with the two contacts in the other sensor's view, with only one way being correct. The probability of confusion may be higher the closer the two ships are together, for instance. Similar confusion can occur with radar contacts in air traffic surveillance, if an attacking fighter jet "hides" very close to a civilian jetliner, which may be their preferred method of approach.

Let the probability of missassociation error in fusion be P^M. When the occasions for inference come in pairs and there is no systematic bias $0 < P^M \leq 1/2$; the upper limit being random guessing. The worst-case guessing condition would presumably occur if the system developer never considered the prospect of two very close occasions for decision/action inferences. To illustrate this condition, another MonteCarlo simulation was run much as before, for $n = 5000$ occasions for inferences in which the two world states 0 and 1 occur with equal relative frequencies, randomly distributed throughout the entire sampling of n occasions. The n occasions for inferences were then grouped into $n/2$ pairs. In one half of those pairs one decision variable was flipped; that is, given the pair of ordered pairs (c_1^A, c_2^A) and (c_1^B, c_2^B) correctly associated, they were replaced by a corresponding missassociated pair (c_1^B, c_2^A) and (c_1^A, c_2^B). In effect, subject A in one source was missassociated with subject matter B in the other in one half of the n occasions for inference. The performance curve for fusion with these missassociations was recomputed and plotted in Fig.(6).

Fusion at this missassociation rate clearly causes more harm than good. Its performance curve lies entirely below and left of each source acting independently (the condition (6)). The fact that missassociation undermines fusion is of course no surprise. And it may look here as if we have set up a straw man, only to knock it down again, deliberately sabotaging fusion in order to illustrate an obvious vulnerability. The intent, however, was to create an analytic frame work in which the vulnerability can be quantified and analyzed, without wishfully ignoring it away, to determine how insidious it might be.

In the present example, if no more than two confusable subjects can be expected, then the fusion performance would presumably fall somewhere between that with perfect association $P^M = 0$, and that when $P^M = 1/2$, depending on the relative frequencies with which unambiguous single occasions and confusable paired occasions for inference appear in the real world. The maximum effective missassociation rate P^M that fusion could support before it would be better not to use it, would be when the fusion performance just equals the performance of the best source, which in this example happens when $P^M \approx 0.20$. In other words, if the system developer has reason to expect confusion rates of $P^M = 0.20$ or higher, it would be better not
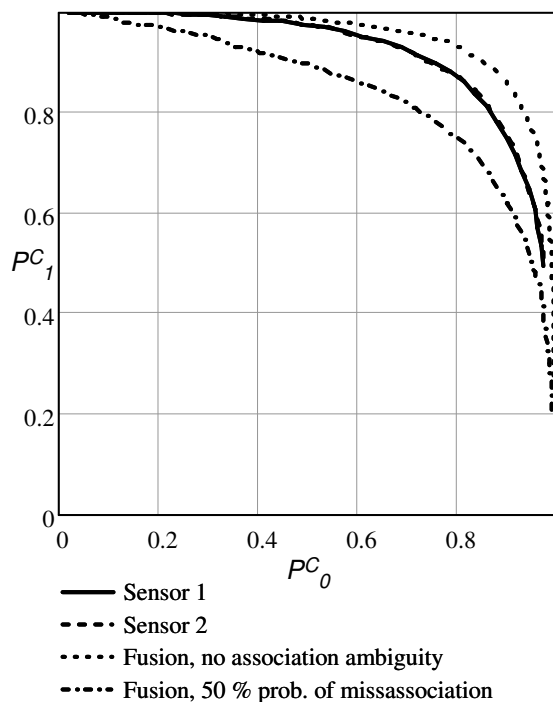
Figure 6: Performance curves for the two sensors working independently (one overlaying the other), and for fusion without and with association errors. Unambiguous fusion outperforms the individual sensors demonstrating the benefit of synthesis. Fusion with missassociation more than cancels those benefits.

to use fusion.

Association typically worsens as the number of confusable occasions for inference increases. For L confusables, $P^M \leq (1 - 1/L)$, driving the performance curves still worse than plotted in Fig.(6). In the heat of battle, of course, the number of confusables may increase dramatically and unexpectedly. The fusion performance will degrade accordingly. These potential degradations must be considered and identified in advance of operations, to be either remedied or addressed with other contingency plans.

4   Conclusions

Association uncertainties are the norm in fusion, especially for complex applications at critical moments. It is only a question of the degree. The treatment given here enters into that question quantitatively.

A generic model of fusion was introduced to quantitatively demonstrate the competing dynamics that inevitable occur to varying degrees in fusion, namely, the increased quality of information due to the combination of independent sources of information, versus the decreased quality owing to association uncertainties. Vulnerability to missassociation has always been recognized, of course, but its impact is modeled quantitatively now within a generic framework that admits further analysis. There are, for instance, perhaps two

ways to reduce association errors: either by 1) gathering and integrating still more information of quite a different nature than that being fused, namely, meta information about situation and context, to drive down the probability of missassociation; or, possibly, by 2) reducing the permutations and combinations of confusables by reducing the number of sources of information used in the fusion. The present analysis may allow experimentation with both options.

Ultimately the analysis would lead (it is hoped) to quantitative engineering rules of thumb that identify good candidates for fusion, and identify situations in which fusion is likely to fail in practice, before the systems are built. These design rules would presumably relate the maximum probability of missassociation P^M to the number and quality of the individual sources of information when fusion is a break-even proposition. One would want to ensure that P^M remained much less than that in critical operations.

References

[1] E.L. Waltz and J.Llinas, *Multisensor data fusion*, Artech House, Norwood, 1990

[2] D.E. Green and J.A. Swets, *Signal detection theory and psychophysics*, Penninsula Publishing, Los Altos California, 1988

[3] B.V. Dasarathy, *Decision Fusion*, IEEE Computer Society Press, Los Alamitos California, 1994

[4] B.V. Dasarathy, "Fusion Strategies for Enhancing Decision Reliability in Multisensor Environments", *Optical Engineering*, Vol. 35, No. 3, pp. 603-616, March 1996.

Document Data Sheet

Security Classification		Project No.
RELEASABLE TO THE PUBLIC		

Document Serial No.	Date of Issue	Total Pages
NURC-PR-2006-003	August 2006	8 pp.

Author(s)

Kessel, Ronald T.

Title

The Dynamics of Information Fusion: Synthetic versus Misassociation.

Abstract

Keywords

Issuing Organization

NATO Undersea Research Centre
Viale San Bartolomeo 400, 19138 La Spezia, Italy

[From N. America:
NATO Undersea Research Centre
(New York) APO AE 09613-5000]

Tel: +39 0187 527 361
Fax:+39 0187 527 700

E-mail: library@nurc.nato.int